

 **PORTAL**  
USPTO

Subscribe (Full Service) [Register \(Limited Service, Free\)](#) [Login](#)  
 Search:  The ACM Digital Library  The Guide  
 "base address" "alternate address" +disk

 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [base address](#) [alternate address](#) [disk](#)

Found 413 of 167,655

Sort results by [relevance](#)   [Save results to a Binder](#)  
 Display results [expanded form](#)   [Search Tips](#)  [Open results in a new window](#)

[Try an Advanced Search](#)  
[Try this search in The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale 

**1 [Implementation aspects of a SPARC V9 complete machine simulator](#)** 

Bill Clarke, Adam Czezowski, Peter Strazdins

January 2002 **Australian Computer Science Communications , Proceedings of the twenty-fifth Australasian conference on Computer science - Volume 4 CRPITS '02**, Volume 24 Issue 1

Publisher: Australian Computer Society, Inc. , IEEE Computer Society Press

Full text available:  [pdf\(1.33 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper we present work in progress in the development of a complete machine simulator for the UltraSPARC, an implementation of the SPARC V9 architecture. The complexity of the UltraSPARC ISA presents many challenges in developing a reliable and yet reasonably efficient implementation of such a simulator. Our implementation includes a heavily object-oriented design for the simulator modules and infrastructure, caching of repeated computations for performance, adding an OS (system call) emu ...

**Keywords:** SMP, SPARC V9 ISA, UltraSPARC, complete machine simulator, execution-driven simulation, object-oriented design

**2 [Disk-directed I/O for MIMD multiprocessors](#)** 

 David Kotz

February 1997 **ACM Transactions on Computer Systems (TOCS)**, Volume 15 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(559.18 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Many scientific applications that run on today's multiprocessors, such as weather forecasting and seismic analysis, are bottlenecked by their file-I/O needs. Even if the multiprocessor is configured with sufficient I/O hardware, the file system software often fails to provide the available bandwidth to the application. Although libraries and enhanced file system interfaces can make a significant improvement, we believe that fundamental changes are needed in the file server software. We prop ...

**Keywords:** MIMD, collective I/O, disk-directed I/O, file caching, parallel I/O, parallel file system

**3**

[Lilith: A personal computer for the software engineer](#) 

N. Wirth

March 1981 **Proceedings of the 5th international conference on Software engineering**

Publisher: IEEE Press

Full text available: [pdf\(1.18 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The personal work station offers significant advantages over the large-scale, central computing facility accessed via a terminal. Among them are availability, reliability, simplicity of operation, and a high bandwidth to the user. Modern technology allows to build systems for high-level language programming with significant computing power for a reasonable price. At the Institut fur Informatik of ETH we have designed and built such a personal computer tailored to the language Mod ...

#### 4 JOYCE: A next generation personal computer



W. R. Franta

September 1980 **Proceedings of the 3rd ACM SIGSMALL symposium and the first SIGPC symposium on Small systems**

Publisher: ACM Press

Full text available: [pdf\(593.06 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper we will detail the JOYCE hardware/software design decisions, cost, status and our plans for continued system development. We examine its features in light of some scheduled applications including CAI (a stand alone PLATO system), database management, electronic mail, etc. We will also detail our experiences relating to the acquisition of components necessary for the realization of JOYCE, including delivery delays, vendor promises and the gap between vendor product ( ...

**Keywords:** Fixed head disk, Microprocessor, Personal computer

#### 5 Application performance and flexibility on exokernel systems



M. Frans Kaashoek, Dawson R. Engler, Gregory R. Ganger, Hector M. Briceño, Russell Hunt, David Mazières, Thomas Pinckney, Robert Grimm, John Jannotti, Kenneth Mackenzie  
October 1997 **ACM SIGOPS Operating Systems Review , Proceedings of the sixteenth ACM symposium on Operating systems principles SOSP '97**, Volume 31 Issue 5

Publisher: ACM Press

Full text available: [pdf\(2.39 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

#### 6 The term retrieval abstract machine



Michael Ley

June 1992 **ACM SIGMOD Record , Proceedings of the 1992 ACM SIGMOD international conference on Management of data SIGMOD '92**, Volume 21 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.22 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Scans through large collections of complex objects often cannot be avoided. Even if sophisticated indexing mechanisms are provided, it may be necessary to evaluate simple predicates against data stored on disk for filtering. For traditional record oriented data models i/o and buffer management are the main bottlenecks for this operation, the interpretation of data structures is straightforward and usually not an important cost factor. For heterogeneously shaped complex objects it may become ...

#### 7 System architectures for computer music

John W. Gordon

◆ June 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 2

Publisher: ACM Press

Full text available:  pdf(4.61 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Computer music is a relatively new field. While a large proportion of the public is aware of computer music in one form or another, there seems to be a need for a better understanding of its capabilities and limitations in terms of synthesis, performance, and recording hardware. This article addresses that need by surveying and discussing the architecture of existing computer music systems. System requirements vary according to what the system will be used for. Common uses for co ...

8 An approach to evaluating time sharing systems: MH-TSS a case study 

◆ Linda S Wright, William A Burnette

January 1976 **ACM SIGMETRICS Performance Evaluation Review**, Volume 5 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.58 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The authors conducted a benchmark measurement of the Murray Hill Time Sharing System (MH-TSS) running on a Honeywell 6000. The object of the test was to duplicate the load normally present on the Murray Hill production system, and measure the system's behavior before and after a major software release and a major hardware improvement. Five different load levels, from 30 to 90 users, were measured for each configuration. This paper discusses the methods used in the design of the experiment and in ...

**Keywords:** event trace, monitoring, operating systems, queuing networks, response time, state transition models

9 Lightweight recoverable virtual memory 

◆ M. Satyanarayanan, Henry H. Mashburn, Puneet Kumar, David C. Steere, James J. Kistler  
December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles SOSP '93**, Volume 27 Issue 5

Publisher: ACM Press

Full text available:  pdf(1.53 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

*Recoverable virtual memory* refers to regions of a virtual address space on which transactional guarantees are offered. This paper describes RVM, an efficient, portable, and easily used implementation of recoverable virtual memory for Unix environments. A unique characteristic of RVM is that it allows independent control over the transactional properties of atomicity, permanence, and serializability. This leads to considerable flexibility in the use of RVM, potentially enlarging the ...

10 The ACTION/FILES file system 

◆ John S. Fisher  
September 1976 **Proceedings of the eighth international conference on APL**

Publisher: ACM Press

Full text available:  pdf(383.56 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

This paper describes some of the aspects of the new implementation of the ACTION/FILES system. It is a part of the ACTION/APL.SV system offered by The Computer Company, Richmond, Virginia. The discussion considers the interface between APL user and the files, as well as the internal structuring of the system database and of the system programs.

11 Lightweight recoverable virtual memory

◆ M. Satyanarayanan, Henry H. Mashburn, Puneet Kumar, David C. Steere, James J. Kistler  
 February 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.73 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Recoverable virtual memory refers to regions of a virtual address space on which transactional guarantees are offered. This article describes RVM, an efficient, portable, and easily used implementation of recoverable virtual memory for Unix environments. A unique characteristic of RVM is that it allows independent control over the transactional properties of atomicity, permanence, and serializability. This leads to considerable flexibility in the use of RVM, potentially enabling ...

**Keywords:** Camelot, Coda, RVM, Unix, logging, paging, persistence, scalability, throughput, truncation

## 12 Simulating the NASA Mass Data Storage Facility

Myron H. MacDougall

June 1974 **Proceedings of the 2nd symposium on Simulation of computer systems**

Publisher: IEEE Press

Full text available:  pdf(736.36 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The NASA/Houston Mass Data Storage Facility (MDSF) is a multi-computer system providing data storage and retrieval services to the Real Time Computer Complex (RTCC) of the NASA Mission Control Center in support of the Skylab program. A large scale simulation model of the MDSF has been developed. There are two major areas of use for this model. One is concerned with determining how to use the present MDSF most effectively in accomplishing its current applications. Here, use of the model incl ...

## 13 Operational characteristics of a hardware-based pattern matcher

Roger L. Haskin, Lee A. Hollaar

March 1983 **ACM Transactions on Database Systems (TODS)**, Volume 8 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.84 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The design and operation of a new class of hardware-based pattern matchers, such as would be used in a backended database processor in a full-text or other retrieval system, is presented. This recognizer is based on a unique implementation technique for finite state automata consisting of partitioning the state table among a number of simple digital machines. It avoids the problems generally associated with implementing finite state machines, such as large state table memories, complex control ...

**Keywords:** backend processors, computer system architecture, finite state automata, full text retrieval systems, text searching

## 14 Shading, surfaces, and collision detection: Mesh mutation in programmable graphics hardware

Le-Jeng Shiue, Vineet Goel, Jorg Peters

July 2003 **Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware**

Publisher: Eurographics Association

Full text available:  pdf(789.95 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

We show how a future graphics processor unit (GPU), enhanced with random read and write to video memory, can represent, refine and adjust complex meshes arising in modeling, simulation and animation. To leverage SIMD parallelism, a general model based on the mesh atlas is developed and a particular implementation without adjacency pointers is proposed in which primal, binary refinement of, possibly mixed, quadrilateral and triangular meshes of arbitrary topological genus, as well as their traversal ...

**15 OBJEKT—a persistent object store with an integrated garbage collector**

 D M Harland, B Beloff  
April 1987 **ACM SIGPLAN Notices**, Volume 22 Issue 4

**Publisher:** ACM Press

Full text available:  [pdf\(1.06 MB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

This paper describes OBJEKT, a single-level persistent storage system designed for the REKURSIV architecture. It will be shown that OBJEKT can be microcoded to implement "objects" efficiently, and that data integrity can be guaranteed by provision of an object oriented instruction set. Particular attention will be paid to its facilities for type and range checking, to its object by-object paging strategy and to ways of enhancing parallelism during garbage collection.

**16 Devirtualizable virtual machines enabling general, single-node, online maintenance**

 David E. Lowell, Yasushi Saito, Eileen J. Samberg  
October 2004 **ACM SIGARCH Computer Architecture News**, **ACM SIGOPS Operating Systems Review**, **ACM SIGPLAN Notices**, **Proceedings of the 11th international conference on Architectural support for programming languages and operating systems ASPLOS-XI**, Volume 32, 38, 39 Issue 5, 5, 11

**Publisher:** ACM Press

Full text available:  [pdf\(174.01 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Maintenance is the dominant source of downtime at high availability sites. Unfortunately, the dominant mechanism for reducing this downtime, cluster rolling upgrade, has two shortcomings that have prevented its broad acceptance. First, cluster-style maintenance over many nodes is typically performed a few nodes at a time, making maintenance slow and often impractical. Second, cluster-style maintenance does not work on single-node systems, despite the fact that their unavailability during maintenance ...

**Keywords:** availability, online maintenance, planned downtime, virtual machines

**17 I/O buffer performance in a virtual memory system**

 Stephen W. Sherman, Richard S. Brice  
July 1976 **ACM SIGSIM Simulation Digest**, **Proceedings of the 4th symposium on Simulation of computer systems ANSS '76**, Volume 7 Issue 4

**Publisher:** IEEE Press, ACM Press

Full text available:  [pdf\(1.06 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this study we construct a simulator of a data base management system running in a virtual memory environment. We use the simulator to investigate the value of using an I/O buffer in this environment. The simulator is driven by trace data obtained with a software probe. The simulator is validated and is used to verify a theoretical model which predicts paging and disk access rates produced by use of an I/O buffer in a virtual memory environment. Results from a multi-factor set of simulations ...

**18 A reliable object-oriented data repository for a distributed computer system**

Liba Svobodova  
December 1981 **Proceedings of the eighth ACM symposium on Operating systems**

**principles****Publisher:** ACM PressFull text available: [pdf\(1.18 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The repository described in this paper is a component of a distributed data storage system for a network of many autonomous machines that might run diverse applications. The repository is a server machine that provides very large, very reliable long-term storage for both private and shared data objects. The repository can handle both very small and very large data objects, and it supports atomic update of groups of objects that might be distributed over several repositories. Each object is ...

**Keywords:** Atomic update, Crash recovery, Distributed data storage system, Memory management, Optical disk, Server, Stable storage

**19 Third Generation Computer Systems**

Peter J. Denning

December 1971 **ACM Computing Surveys (CSUR)**, Volume 3 Issue 4**Publisher:** ACM PressFull text available: [pdf\(3.52 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The common features of third generation operating systems are surveyed from a general view, with emphasis on the common abstractions that constitute at least the basis for a "theory" of operating systems. Properties of specific systems are not discussed except where examples are useful. The technical aspects of issues and concepts are stressed, the nontechnical aspects mentioned only briefly. A perfunctory knowledge of third generation systems is presumed.

**20 A comparative study of environments for database system implementation**

M. Teresa Suarez Fernandez, H. Rex Hartson

August 1978 **ACM SIGMINI Newsletter, Proceedings of the first SIGMINI symposium on Small systems SIGMINI '78**, Volume 4 Issue 4**Publisher:** ACM PressFull text available: [pdf\(742.67 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Two specific environments for implementation of experimental DataBase Management Systems (DBMS) are compared, especially in light of performance measurement: IBM 370 and HP-2100A. The HP-2100A was selected for a skeletal implementation to study performance parameters further. The DBMS structure, some performance results, and several difficulties encountered are described. The HP-2100A minicomputer did provide "hands-on" experience and more direct control to customize disk access ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)

## Refine Search

## Search Results -

Terms	Documents
L7 and L8	76

US Pre-Grant Publication Full-Text Database  
US Patents Full-Text Database  
US OCR Full-Text Database  
EPO Abstracts Database  
JPO Abstracts Database  
Derwent World Patents Index  
IBM Technical Disclosure Bulletins

**Database:** EPO Abstracts Database  
JPO Abstracts Database  
Derwent World Patents Index  
IBM Technical Disclosure Bulletins

Search:	<input type="text"/>	<input type="button" value="Refine Search"/>
<input type="button" value="Recall Text"/>		<input type="button" value="Clear"/>
<input type="button" value="Interrupt"/>		

## Search History

DATE: Thursday, December 08, 2005 [Printable Copy](#) [Create Case](#)

<u>Set</u>	<u>Hit</u>	
<u>Name</u>	<u>Count</u>	<u>Set</u>
side by side	result set	name
<i>DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=NO; OP=ADJ</i>		
<u>L9</u> 17 and L8	76	<u>L9</u>
<u>L8</u> 14 same 13	2185	<u>L8</u>
<u>L7</u> 12 and 13 and 15 and L6	278	<u>L7</u>
<u>L6</u> disk	908226	<u>L6</u>
<u>L5</u> compar\$5 with address\$2	71905	<u>L5</u>
<u>L4</u> (designat\$4 or redesignat\$4 or new or another or alternat\$5) near3 address\$2	79759	<u>L4</u>
<u>L3</u> (base or original) near2 (address\$3 or location)	44869	<u>L3</u>
<u>L2</u> alternat\$4 near2 (address\$3 or location)	15501	<u>L2</u>
<u>L1</u> base adj address\$2	9627	<u>L1</u>

END OF SEARCH HISTORY